

Characterizing timing and memory-requirements of the \mathcal{F} -statistic implementations in LALSuite

Reinhard Prix*

Max-Planck-Institut für Gravitationsphysik, Albert-Einstein-Institut, D-30167 Hannover, Germany

(Dated: 2017-07-18 17:50:43 +0200; commitID: 1c97c53-CLEAN; LIGO-T1600531-v4+)

We develop a generic (method-independent) timing model for the \mathcal{F} -statistic computation. We present and quantify method-specific timing and memory models for the Resampling-FFT method, as well as the (older) Demod method. The timing models are found to agree well with the measured CPU times with a standard deviation on the error of less than 10%.

I. THE \mathcal{F} -STATISTIC TIMING MODEL

In the following we denote as \mathcal{T} any times (per detector) referring to a single call to `XLALComputeFstat()`, which computes the \mathcal{F} -statistic over a vector of N_{Fbin} frequency bins, τ any times per detector per output \mathcal{F} -statistic frequency bin N_{Fbin} , and $\tau^{(0)}$ any “fundamental” timing coefficients that don’t scale with any of the algorithm parameters (at least approximately), and which should only depend on the hardware and code optimization settings.

The general \mathcal{F} -statistic timing model can be expressed in terms of two contributions:

1. the average *core* time $\tau_{\mathcal{F}}^{\text{core}}$ to compute a single \mathcal{F} -statistic frequency bin for one detector, excluding the time to compute any quantities that are stored in the buffer, and
2. the average time $\tau_{\mathcal{F}}^{\text{buffer}}$ per frequency bin per detector to (re-)compute all buffered quantities

We can write the *effective* average \mathcal{F} -statistic time $\tau_{\mathcal{F}}^{\text{eff}}$ per template per detector as

$$\tau_{\mathcal{F}}^{\text{eff}} = \tau_{\mathcal{F}}^{\text{core}} + b \tau_{\mathcal{F}}^{\text{buffer}}, \quad (1)$$

where the “buffer miss fraction” $b \in [0, 1]$ quantifies how many times the buffer needed to be recomputed per call to `XLALComputeFstat()` per output frequency bin, i.e. $b = 0$ is 100% efficient buffering, i.e. the buffer never needed to be re-computed, while $b = 1$ means the buffer needed to be re-computed for every call to `XLALComputeFstat()`. The buffer contains sky-dependent quantities (antenna-pattern functions, SSB time delays, SSB-resampled time-series, ...) as well as quantities that depend on the binary-orbital parameters (such as binary-orbital time delays and corresponding resampled time-series). Typically we can re-use the buffer for all spindowns for given sky-position and binary-orbital parameters, and so in most cases we will have

$$b = \frac{1}{N_{\{\dot{f}, \ddot{f}, \dots\}}}, \quad (2)$$

in terms of the number of spindown templates of all orders computed per sky-position and binary-orbital parameters.

The timing collection within the \mathcal{F} -statistic code measures and outputs these quantities as

$$\text{tauF_eff} \equiv \tau_{\mathcal{F}}^{\text{eff}} = \tau_{\mathcal{F}}^{\text{core}} + b \tau_{\mathcal{F}}^{\text{buffer}}, \quad (3)$$

$$\text{tauF_core} \equiv \tau_{\mathcal{F}}^{\text{core}}, \quad (4)$$

$$\text{tauF_buffer} \equiv \tau_{\mathcal{F}}^{\text{buffer}}, \quad (5)$$

$$\text{bufferMissFrac} \equiv b. \quad (6)$$

A *timing model* consists of expressing the \mathcal{F} -statistic times $\tau_{\mathcal{F}}^{\text{core}}$ and $\tau_{\mathcal{F}}^{\text{buffer}}$ in terms of algorithm parameters and fundamental timing coefficients $\tau_{\dots}^{(0)}$, which allows us to predict \mathcal{F} -statistic compute times as a function of algorithm parameters, assuming we have measured the fundamental timing coefficients for a given code and hardware.

* reinhard.prix@aei.mpg.de

A. Resampling timing model

1. Notation

There are essentially four different vector lengths over which various operations are performed in the Resampling-FFT \mathcal{F} -statistic code:

N_{Fbin} : The user-requested number of output frequency bins, with spacing df , for which the \mathcal{F} -statistic is computed.

$N_{\text{samp}}^{\text{DET}}$: The (maximum over detectors) number of samples in the input timeseries (in the detector frame).

$N_{\text{samp}}^{\text{SRC}}$: The (maximum over detectors) number of samples of the input timeseries, interpolated into the source-frame.

$N_{\text{samp}}^{\text{FFT}}$: The length of the (zero-padded) timeseries on which the FFT is performed.

which are given by the following relations:

The sampling step dt_{DET} of the input timeseries in the detector frame (DET) is given in terms of the frequency bandwidth Δf_{sft} of the input SFTs, namely

$$dt_{\text{DET}} \equiv \frac{1}{\Delta f_{\text{sft}}}, \quad (7)$$

and the number of detector-frame time-samples are therefore

$$N_{\text{samp}}^{\text{DET}} = \frac{T_{\text{coh}}}{dt_{\text{DET}}} = T_{\text{coh}} \Delta f_{\text{sft}}, \quad (8)$$

where T_{coh} is the (maximum over IFOs) maximum coherent data span, which is determined by the actual SFTs used and can therefore generally be less than the user-requested (e.g. via segment list) span. The duration T_{FFT} of the final timeseries to be FFT'ed is determined by the user-requested frequency resolution df , namely

$$T_{\text{FFT}} = \frac{D}{df}, \quad \text{with } D \equiv \lceil T_{\text{coh}} df \rceil, \quad (9)$$

where $D \in \mathbb{N}^+$ ensures that we always have $T_{\text{FFT}} \geq T_{\text{coh}}$. A choice $D > 1$ is required when the user asked for a coarse frequency resolution $df > 1/T_{\text{coh}}$ in order to still use all the data (i.e. we always use zero-padding rather than truncating the input timeseries) and therefore not lose SNR. In the regular case of a frequency resolution $df < 1/T_{\text{coh}}$ we have $D = 1$. The *decimation* factor D amounts to internally using a finer frequency resolution

$$df_{\text{internal}} \equiv \frac{df}{D} \leq \frac{1}{T_{\text{coh}}}, \quad (10)$$

and then returning only the user-requested frequency bins with resolution df , namely by returning only every D 'th frequency bin.

With the original detector-frame sampling step dt_{DET} this would correspond to a number of time samples to be FFT'ed:

$$N_{\text{samp}}^{\text{FFT}(0)} = \frac{T_{\text{FFT}}}{dt_{\text{DET}}} = \frac{\Delta f_{\text{sft}}}{df_{\text{internal}}}. \quad (11)$$

However, in order to ensure the most efficient (and most consistent) FFT performance, we round the number of FFT bins up to the next power of 2, namely

$$N_{\text{samp}}^{\text{FFT}} = 2^{\lceil \log_2 N_{\text{samp}}^{\text{FFT}(0)} \rceil}, \quad (12)$$

which is achieved by decreasing the sampling step (i.e. increasing the frequency band) of the resampled timeseries in the source frame (SRC), i.e.

$$dt_{\text{SRC}} \equiv \frac{T_{\text{FFT}}}{N_{\text{samp}}^{\text{FFT}}}. \quad (13)$$

The (longest) coherent detector-frame time-series of span T_{coh} interpolated into the source-frame will therefore have $N_{\text{samp}}^{\text{SRC}}$ samples, given as

$$N_{\text{samp}}^{\text{SRC}} \equiv \frac{T_{\text{coh}}}{dt_{\text{SRC}}} = \mathcal{R} N_{\text{samp}}^{\text{FFT}}, \quad (14)$$

where we defined the (internal) frequency resolution $\mathcal{R} \leq 1$ in “natural” units $1/T_{\text{coh}}$ as

$$df_{\text{internal}} = \mathcal{R} \frac{1}{T_{\text{coh}}}, \quad (15)$$

namely

$$\mathcal{R} \equiv \frac{T_{\text{coh}}}{T_{\text{FFT}}} = T_{\text{coh}} df_{\text{internal}} \leq 1. \quad (16)$$

From Eq. (8) and Eq. (11) we find a similar relation to Eq. (14) for $N_{\text{samp}}^{\text{DET}}$, namely

$$N_{\text{samp}}^{\text{DET}} = \mathcal{R} N_{\text{samp}}^{\text{FFT}(0)}. \quad (17)$$

B. Required SFT frequency band Δf_{sft}

When the caller requests a physical SRC-frame frequency band Δf to be analyzed, we need to expand this to Δf_{load} when loading the data from SFTs: this wider band contains an extra “frequency drift” sideband Δf_{drift} to account for Doppler shifts and spindowns (computed from `XLALCWSignalCoveringBand()` [1]), and an additional 16 extra SFT bins (8 on either side¹) added to use similar noise bins as Demod near the SFT edges, i.e.

$$\Delta f_{\text{load}} = \Delta f + \Delta f_{\text{drift}} + \frac{16}{T_{\text{sft}}}, \quad (18)$$

$$\Delta f_{\text{drift}} \approx 2.12 \times 10^{-4} \left(f_{\text{max}} + |\dot{f}|_{\text{max}} \frac{T_{\text{span}}}{2} + |\ddot{f}|_{\text{max}} \frac{T_{\text{span}}^2}{8} + \dots \right), \quad (19)$$

where T_{span} denotes the total duration of data included in this calculation. The handling of this can differ between semi-coherent search codes, for example for the GCT code, T_{span} denotes the *total* span of data over all semi-coherent segments, whereas for the Weave code, this could be computed (and data loaded) on a per-segment basis, such that for Weave $T_{\text{span}} = T_{\text{coh}}$. An additional “transition band” is added internally in order to allow for the roll-off from the Hamming-windowed sinc-interpolation used in Barycentric resampling, which further increases this to

$$\Delta f_{\text{sft}} = \Delta f_{\text{load}} \left(1 + \frac{4}{2\text{Dterms} + 1} \right), \quad (20)$$

where `Dterms` is the (user-specified) number of sinc-kernel terms to use (on either side, therefore the window-size is $2\text{Dterms} + 1$) in the barycentric resampling interpolation.

Note that these expressions (especially Δf_{drift}) are only approximations of what is computed in the F-stat codes, and due to the power-of-two rounding of Eq. (12), this can in some cases lead to relatively large deviations from the actual measured runtimes. For the most accurate and reliable prediction use the `octapps` function `predictFstatTimeAndMemory()` [2].

The “physical” number of maximal output frequency bins is

$$N_{\text{Fbin}}^{\text{max}} = \frac{\Delta f}{df}, \quad (21)$$

while calls to `XLALComputeFstat()` can request to compute fewer output frequency bins $N_{\text{Fbin}} \leq N_{\text{Fbin}}^{\text{max}}$, provided they fall into the requested frequency band Δf . Note that for semi-coherent codes such as the GCT code, Δf will not be just the user-requested search frequency band Δf_0 , but will typically include additional sideband bins, for example in the case of the GCT code we’ll have a wider band of

$$\Delta f = \Delta f_0 + \frac{1}{2} \left(T_{\text{span}} d\dot{f} + T_{\text{span}}^2 d\ddot{f} \right), \quad (22)$$

where the extra sideband is referred to as `extraBinsFstat` in the GCT code.

¹ called `extraBinsMethod` in the code

1. Timing model contributions

We can break the time to compute the Resampling \mathcal{F} -statistic into the following contributions: the *core* compute time $\mathcal{T}_{\mathcal{F}}^{\text{core}}$ has the following contributions $\mathcal{T}_{\mathcal{F}}^{\text{core}} = \mathcal{T}_{\text{Fbin}} + \mathcal{T}_{\text{spin}} + \mathcal{T}_{\text{FFT}}$, which scale differently with search parameters, namely

1. Apply spindown-correction and frequency shift to source-frame timeseries, which scales with $N_{\text{samp}}^{\text{SRC}}$, and therefore

$$\mathcal{T}_{\text{spin}} = N_{\text{samp}}^{\text{SRC}} \tau_{\text{spin}}^{(0)} = \mathcal{R} N_{\text{samp}}^{\text{FFT}} \tau_{\text{spin}}^{(0)}. \quad (23)$$

2. Apply the FFT, which operates on the zero-padded timeseries of length $N_{\text{samp}}^{\text{FFT}}$, therefore

$$\mathcal{T}_{\text{FFT}} = 5 N_{\text{samp}}^{\text{FFT}} \log_2(N_{\text{samp}}^{\text{FFT}}) \tau_{\text{FFT}}^{(0)}, \quad (24)$$

which was defined in such a way that we can directly the FFT timing coefficient $\tau_{\text{FFT}}^{(0)}$ to the benchmark “speed” definition of FFTW in <http://www.fftw.org/speed/>. This is defined in terms of “mega-flops” as $\text{mflops} \equiv 5 N_{\text{samp}}^{\text{FFT}} \log_2(N_{\text{samp}}^{\text{FFT}}) / (\mathcal{T}_{\text{FFT}} 10^6)$, and so FFTW’s “flops” are simply

$$\text{FFTwflops} = \frac{1}{\tau_{\text{FFT}}^{(0)}}. \quad (25)$$

The FFT timing coefficient $\tau_{\text{FFT}}^{(0)}$ is found to be relatively constant for FFTs of length $N_{\text{samp}}^{\text{FFT}} > 2^{18} = 262144$, while it is noticeably smaller but also more variable for shorter FFTs (in agreement with FFTW’s benchmark results for single-precision complex FFTs on 64bit hardware shown here: <http://www.fftw.org/speed/CoreDuo-3.0GHz-icc64/>). In practice we therefore use two FFT timing coefficients, $\tau_{\text{FFT}}^{(0)} (N \leq 2^{18})$ and $\tau_{\text{FFT}}^{(0)} (N > 2^{18})$.

3. various operations on the N_{Fbin} output bins, such as copying the bins, normalizing them, computing \mathcal{F} from F_a, F_b and summing them over detectors $F_{a,b} = \sum_X F_{a,b}^X$ (here considering time per detector), which can be summarized as a time contribution per output frequency bin $\tau_{\text{Fbin}}^{(0)}$, internally composed of these parts

$$\mathcal{T}_{\text{Fbin}} = \mathcal{T}_{\text{copy}} + \mathcal{T}_{\text{norm}} + \mathcal{T}_{\text{sumFabX}} + \mathcal{T}_{\text{Fab2F}} = N_{\text{Fbin}} \tau_{\text{Fbin}}^{(0)}. \quad (26)$$

We therefore obtain the core resampling time per detector per output frequency bin N_{Fbin} as

$$\tau_{\mathcal{F}}^{\text{core}} \equiv \frac{\mathcal{T}_{\mathcal{F}}^{\text{core}}}{N_{\text{Fbin}}} = \tau_{\text{Fbin}}^{(0)} + \frac{N_{\text{samp}}^{\text{FFT}}}{N_{\text{Fbin}}} \left[\mathcal{R} \tau_{\text{spin}}^{(0)} + 5 \log_2(N_{\text{samp}}^{\text{FFT}}) \tau_{\text{FFT}}^{(0)} \right]. \quad (27)$$

We also need to compute all the buffered quantities: antenna-patterns and source-frame time-delays, barycenter and interpolate the input detector-frame time-series into the source frame. This contribution contains some terms that scale with the number of input SFTs, and others that scale with the number of input samples of the SFTs (in case of gaps), as well as $N_{\text{samp}}^{\text{SRC}}$ and $N_{\text{samp}}^{\text{DET}}$. However, we postulate the following approximate scaling relation for the dominant contribution (barycentering) as

$$\mathcal{T}_{\mathcal{F}}^{\text{buffer}} \sim \mathcal{T}_{\text{bary}} \sim N_{\text{samp}}^{\text{SRC}} \tau_{\text{bary}}^{(0)} = \mathcal{R} N_{\text{samp}}^{\text{FFT}} \tau_{\text{bary}}^{(0)}, \quad (28)$$

and so

$$\tau_{\mathcal{F}}^{\text{buffer}} \equiv \frac{\mathcal{T}_{\mathcal{F}}^{\text{buffer}}}{N_{\text{Fbin}}} = \mathcal{R} \frac{N_{\text{samp}}^{\text{FFT}}}{N_{\text{Fbin}}} \tau_{\text{bary}}^{(0)}. \quad (29)$$

The complete Resampling timing model for the effective average time per template per detector of Eq. (1) can therefore be expressed as

$$\tau_{\mathcal{F}}^{\text{eff}} = \tau_{\text{Fbin}}^{(0)} + \frac{N_{\text{samp}}^{\text{FFT}}}{N_{\text{Fbin}}} \left[\mathcal{R} \left(\tau_{\text{spin}}^{(0)} + b \tau_{\text{bary}}^{(0)} \right) + 5 \log_2(N_{\text{samp}}^{\text{FFT}}) \tau_{\text{FFT}}^{(0)} \right]. \quad (30)$$

C. Demod Timing Model

The timing model for the Demod implementation of the \mathcal{F} -statistic is much simpler. The core Demod \mathcal{F} -statistic time $\tau_{\mathcal{F}}^{\text{core}}$ per detector is simply proportional to the number of SFTs N_{sft} from the single detector, i.e.

$$\tau_{\mathcal{F}}^{\text{core}} = N_{\text{sft}} \tau_{\text{core,LD}}^{(0)}, \quad (31)$$

where $\tau_{\text{core,LD}}^{(0)}$ is the fundamental timing coefficient expressing core Demod time *per SFT*.

In the Demod case, the time $\tau_{\mathcal{F}}^{\text{buffer}}$ to compute all the buffered quantities (sky-dependent antenna-patterns and sky- and binary-orbital dependent timings) for one detector is simply proportional to the number of SFTs from that detector, and we can obtain

$$\tau_{\mathcal{F}}^{\text{buffer}} = \frac{N_{\text{sft}}}{N_{\text{Fbin}}} \tau_{\text{buffer,LD}}^{(0)}, \quad (32)$$

in terms of a fundamental timing coefficient $\tau_{\text{buffer,LD}}^{(0)}$, which denotes the time to compute all the buffered quantities for a single SFT. We can therefore write the full Demod timing model $\tau_{\mathcal{F}}^{\text{eff}}$ for the time per template per detector as

$$\tau_{\mathcal{F}}^{\text{eff}} = N_{\text{sft}} \left[\tau_{\text{core,LD}}^{(0)} + \frac{b}{N_{\text{Fbin}}} \tau_{\text{buffer,LD}}^{(0)} \right], \quad (33)$$

in terms of the number of SFTs N_{sft} coherently analyzed from one detector.

II. MEMORY MODEL

In a similar way to the timing model we can enumerate the amount of memory required to perform the calculation of the \mathcal{F} -statistic.

A. Resampling memory model

- **Objects stored for each detector and each segment of a semi-coherent search:**

- the original SFTs turned into a detector-frame COMPLEX8 timeseries:

$$\text{mem}[\text{COMPLEX8TimeSeries-DET}] = N_{\text{samp}}^{\text{DET}} \times \text{mem}[\text{C8}].$$

- two timeseries interpolated and barycentered into the source frame, multiplied by $a(t)$ or $b(t)$, respectively:

$$\text{mem}[\text{COMPLEX8TimeSeries-SRC-[a|b]}] = 2 N_{\text{samp}}^{\text{SRC}} \times \text{mem}[\text{C8}].$$

- the FFTW plan: we’re *assuming* it should be roughly $N_{\text{samp}}^{\text{FFT}}$ COMPLEX8 numbers(?)

$$\text{mem}[\text{FFT-plan}] \approx N_{\text{samp}}^{\text{FFT}} \times \text{mem}[\text{C8}].$$

- **Objects stored only once for all segments, using a shared “workspace”:**

- Two COMPLEX8 vectors for temporary storage of SRC-frame timeseries:

$$\text{mem}[\text{TStmp[1|2]-SRC}] = 2 N_{\text{samp}}^{\text{SRC}} \times \text{mem}[\text{C8}],$$

- One REAL8 timeseries holding time-differences between SRC and DET frames (only used for barycentering:

$$\text{mem}[\text{SRCtimes-DET}] = N_{\text{samp}}^{\text{SRC}} \times \text{mem}[\text{R8}].$$

- the input and output vectors for the FFT (currently *not* using in-place transform):

$$\text{mem}[\text{TS-FFT + FabX-Raw}] = 2 N_{\text{samp}}^{\text{FFT}} \times \text{mem}[\text{C8}],$$

– temporary storage of $F_{a,b}^X$ and $F_{a,b}$:

$$\text{mem}[\text{Fab}] = 4 N_{\text{Fbin}} \times \text{mem}[\text{C8}] .$$

where $\text{mem}[\text{C8}] = \text{mem}[\text{R8}] = 8$ bytes and $\text{mem}[\text{R4}] = 4$ bytes.

We can therefore express the different memory blocks as

$$\text{mem}[\text{ResampDataPerDetPerSeg}] = \left[\mathcal{R} N_{\text{samp}}^{\text{FFT}(0)} + (1 + 2\mathcal{R}) N_{\text{samp}}^{\text{FFT}} \right] \times 8 \text{ bytes} , \quad (34)$$

$$\text{mem}[\text{ResampWorkspace}] = \left[(2 + 3\mathcal{R}) N_{\text{samp}}^{\text{FFT}} + 4N_{\text{Fbin}} \right] \times 8 \text{ bytes} , \quad (35)$$

$$(36)$$

and the total (multi-segment, multi-detector) “internal” memory usage of the Resampling \mathcal{F} -statistic calculation is therefore

$$\text{mem}[\text{Resamp}] = N_{\text{seg}} N_{\text{det}} \times \text{mem}[\text{ResampDataPerDetPerSeg}] + \text{mem}[\text{ResampWorkspace}] . \quad (37)$$

Note that this accounting does not including the “external” memory required to store the returned \mathcal{F} -statistic results from each call to `XLALComputeFstat()`, namely

$$\text{mem}[\text{F-stats return}] = N_{\text{Fbin}} \times \text{mem}[\text{R4}] ,$$

$$\text{mem}[\text{per-IFO } \{\mathcal{F}^X\}] = N_{\text{det}} N_{\text{Fbin}} \times \text{mem}[\text{R4}] ,$$

$$\text{mem}[\text{per-segment per-IFO } \{\mathcal{F}^{X,\ell}\}] = N_{\text{seg}} N_{\text{det}} N_{\text{Fbin}} \times \text{mem}[\text{R4}] ,$$

depending on whether the user requested just return of \mathcal{F} -statistic, or per-detector \mathcal{F}^X , or also per-segment per-detector $\mathcal{F}^{X,\ell}$. This memory is handled by the calling code and has to be accounted for according to its usage of this memory (ie whether it is stored, or added to a toplist immediately, etc).

B. Demod Memory Model

The Demod code needs essentially no internal workspace other than the memory to hold the input SFTs, i.e.

$$\text{mem}[\text{DemodDataPerSFT}] = (\Delta f_{\text{load}} T_{\text{sft}}) \times 8 \text{ bytes} , \quad (38)$$

and so the total (multi-segment, multi-detector) Demod memory “internal” usage is

$$\text{mem}[\text{Demod}] = N_{\text{sft}}^{\text{all}} \times \text{mem}[\text{DemodDataPerSFT}] , \quad (39)$$

where $N_{\text{sft}}^{\text{all}}$ is the total number of all SFTs over all segments and detectors, i.e. $N_{\text{sft}}^{\text{all}} = N_{\text{det}} N_{\text{seg}} N_{\text{sft}}$, if N_{sft} is the average number of SFTs per detector per segment.

III. TESTING THE TIMING MODEL

A. Timing on a Lenovo Thinkpad T520

CPU: Intel(R) Core(TM) i7-2620M CPU, 2.70GHz,

L2 cache: 4MB

Extensions: sse sse2 sse3 sse4_1 sse4_2 avx

Kernel: Linux 3.16.0-4-amd64

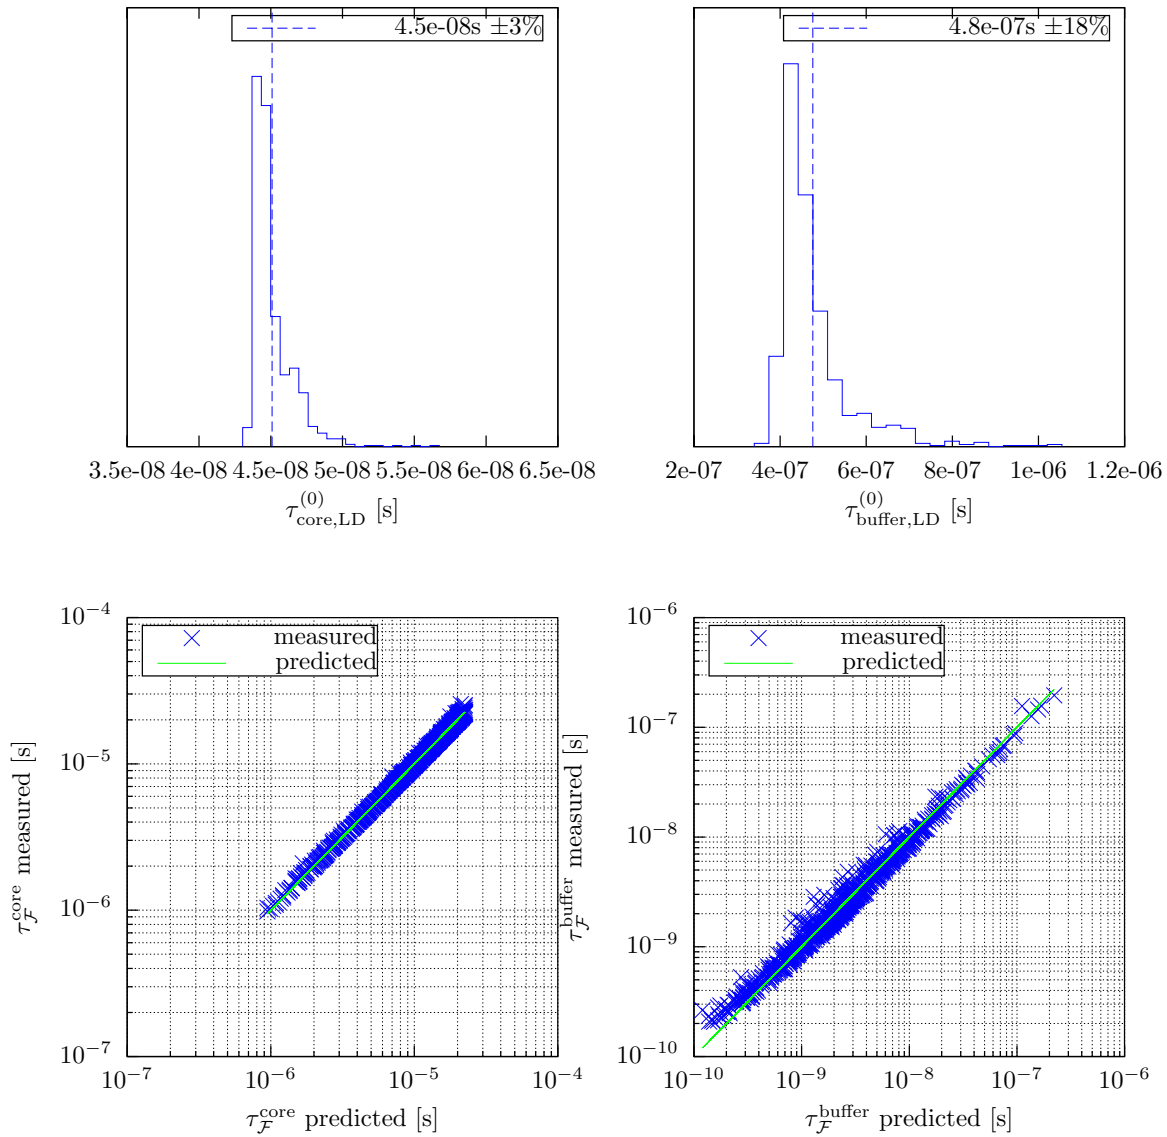
Compiler: Debian clang version 3.5.0-10 (tags/RELEASE_350/final) (based on LLVM 3.5.0)

The timing measurements use `lalapps_ComputeFstatBenchmark`:

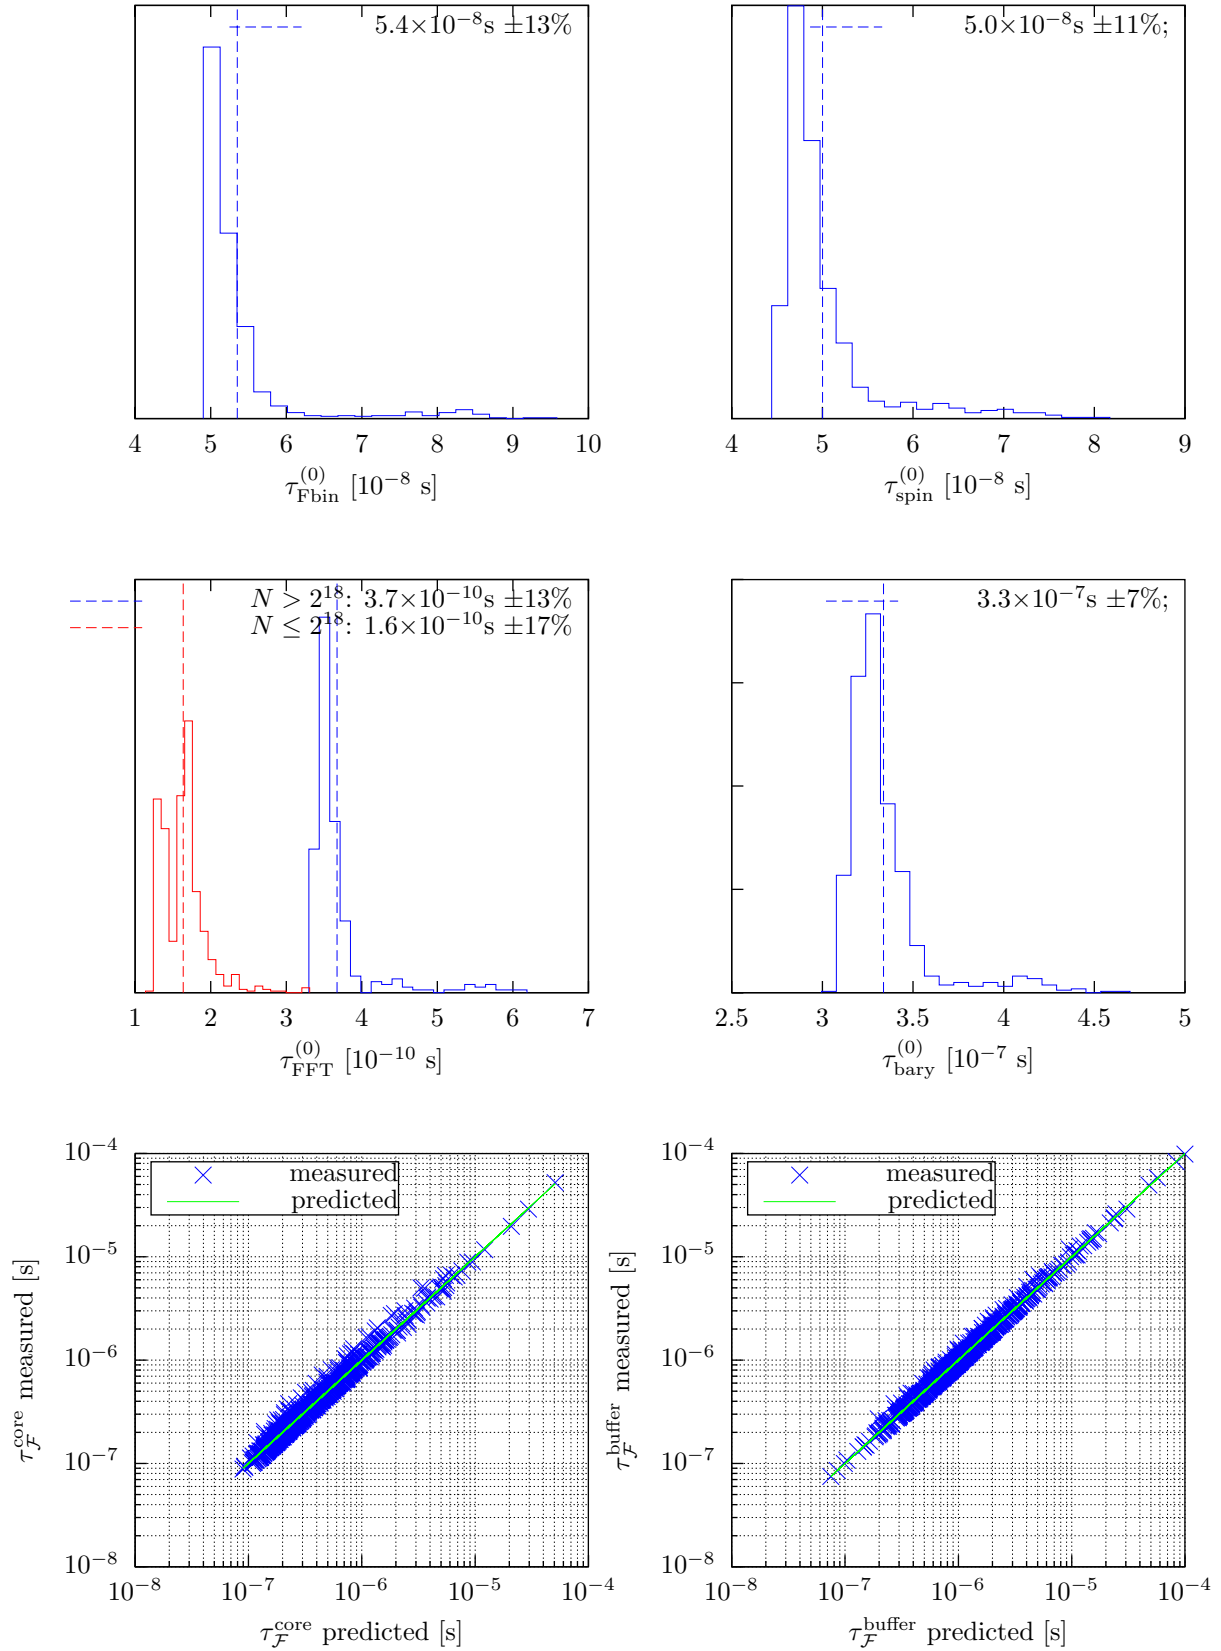
```
$ lalapps_ComputeFstatBenchmark --FstatMethod="ResampBest" --numSegments=1 --numTrials=1000
```

for Resampling, and `--FstatMethod="DemodBest"` for Demod. The model comparison was done using the `octapps` function `predictFstatTimeAndMemory()`.

1. Demod timing: Thinkpad T520



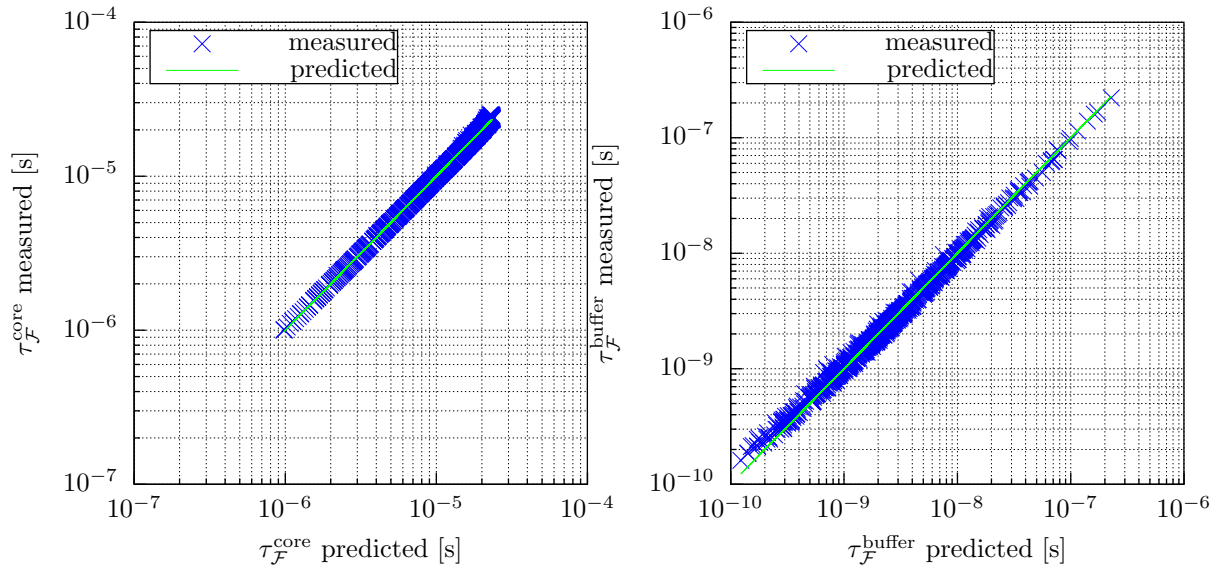

2. Resampling timing: Thinkpad T520



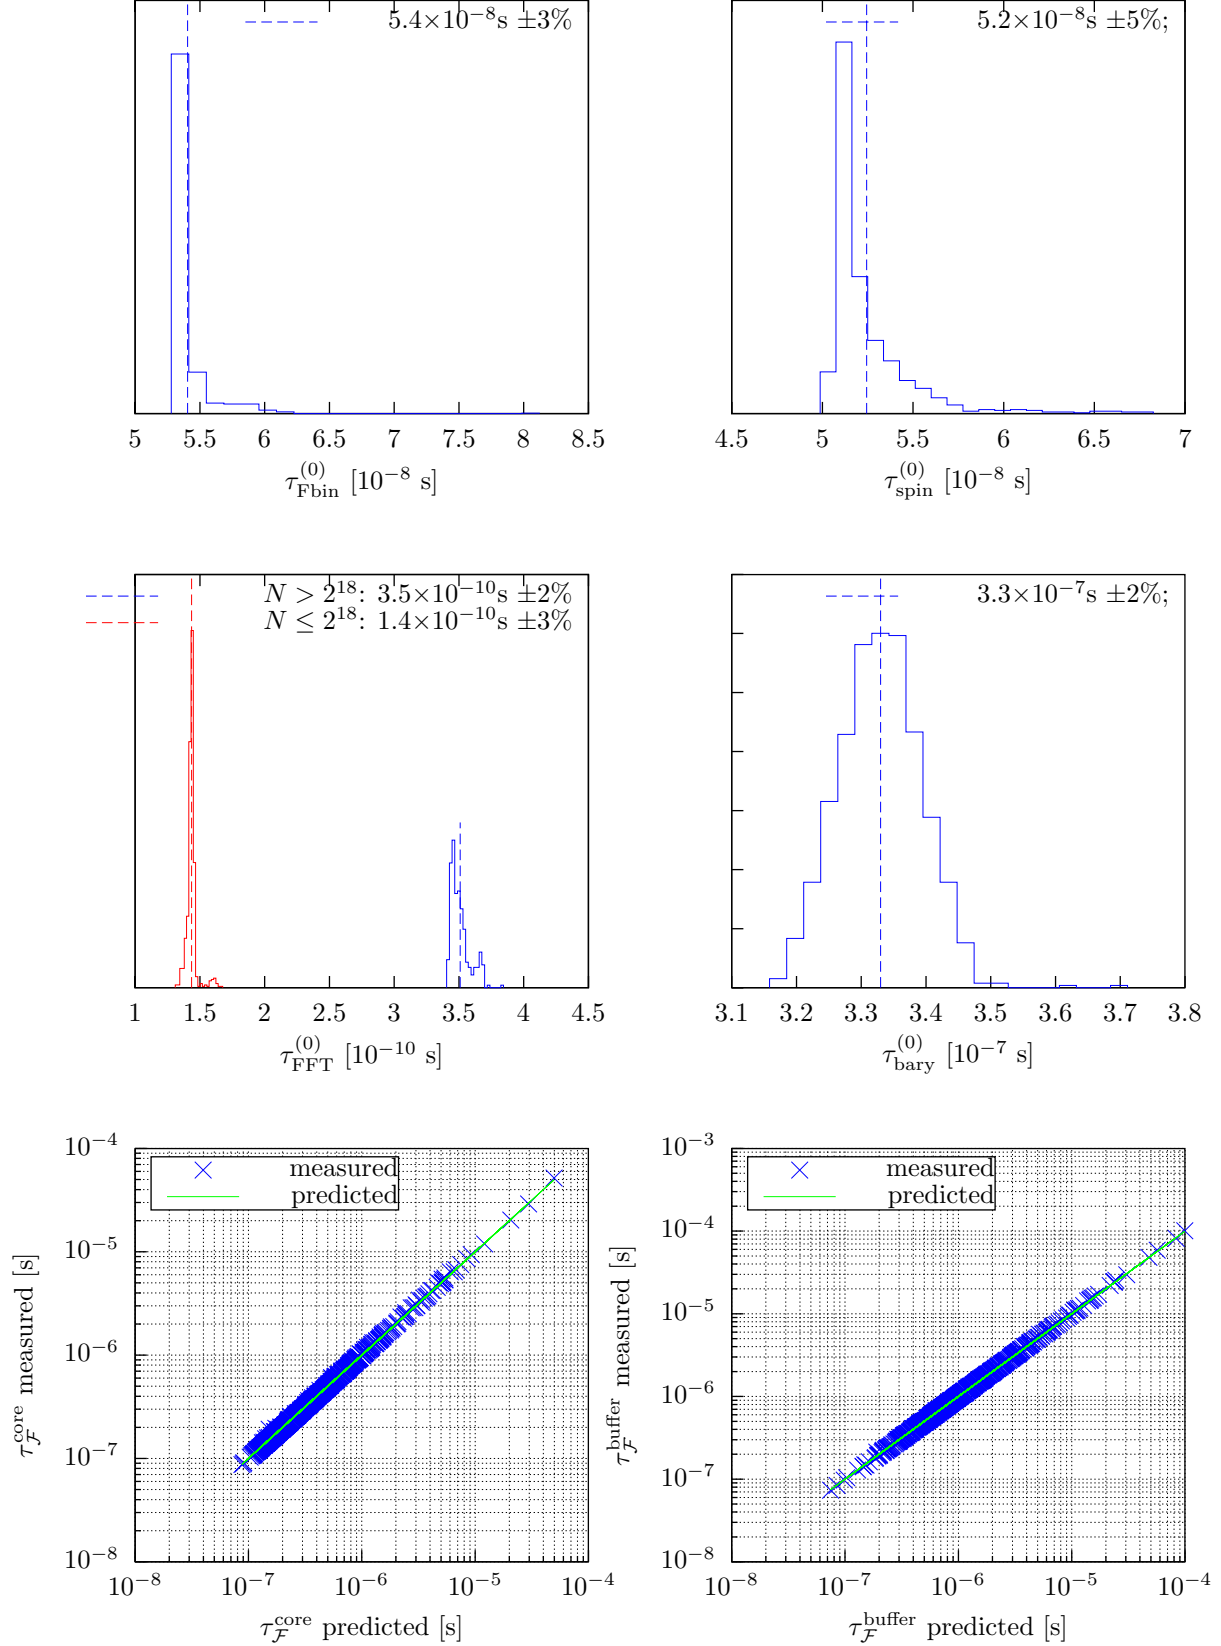
B. Timing on ATLASDEV1

CPU: Intel(R) Xeon(R) CPU E5-1650 0 @ 3.20GHz
 L2 cache: 12MB
 Extensions: sse sse2 sse3 sse4_1 sse4_2 avx
 Kernel: Linux 3.16.0-4-amd64
 Compiler: gcc (Debian 4.9.2-10) 4.9.2

1. Demod timing: ATLASDEV1



2. Resampling timing: ATLASDEV1



C. Average timing coefficients for E@H

As a reasonable “guess” for E@H timings, one could use the average timing coefficients. The following table summarizes the individual and average timing coefficients, all times are in seconds, and the FFT timing coefficients are for $N_{\text{samp}}^{\text{FFT}} \leq 2^{18}$ (“lo”) and $N_{\text{samp}}^{\text{FFT}} > 2^{18}$ (“hi”) respectively:

Hardware	Demod		Resampling			
	$\tau_{\text{core,LD}}^{(0)}$	$\tau_{\text{bary,LD}}^{(0)}$	$\tau_{\text{Fbin}}^{(0)}$	$\tau_{\text{spin}}^{(0)}$	$\tau_{\text{FFT}}^{(0)}$ [lo,hi]	$\tau_{\text{bary}}^{(0)}$
T520-new	4.5×10^{-8}	4.8×10^{-7}	5.4×10^{-8}	5.0×10^{-8}	$[1.6 \times 10^{-10}, 3.7 \times 10^{-10}]$	3.3×10^{-7}
atlasdev1-new	4.7×10^{-8}	4.9×10^{-7}	5.4×10^{-8}	5.2×10^{-8}	$[1.4 \times 10^{-10}, 3.5 \times 10^{-10}]$	3.3×10^{-7}
Average	4.6×10^{-8}	4.8×10^{-7}	5.4×10^{-8}	5.1×10^{-8}	$[1.5 \times 10^{-10}, 3.6 \times 10^{-10}]$	3.3×10^{-7}

[1] *LSC Algorithm Library - LALSuite*, FreeSoftware (GPL), URL <https://wiki.ligo.org/DASWG/LALSuite>.

[2] *Octapps - an LSC Octave library*, FreeSoftware (GPL), URL <https://gitlab.aei.uni-hannover.de/octapps/octapps>.